

03-03-00

A



UTILITY PATENT APPLICATION TRANSMITTAL (Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
PO9-99-092 (12865)

Total Pages in this Submission
59 pages

TO THE ASSISTANT COMMISSIONER FOR PATENTS

Box Patent Application
Washington, D.C. 20231

Transmitted herewith for filing under 35 U.S.C. 111(a) and 37 C.F.R. 1.53(b) is a new utility patent application for an invention entitled:

INTELLIGENT WORKSTATION SIMULATION - SIMULATION AT PROTOCOL STACK LEVEL 2

and invented by:

Max M. Maurer

1C530 U.S. PTO
09/516708
03/02/00

If a **CONTINUATION APPLICATION**, check appropriate box and supply the requisite information:

Continuation Divisional Continuation-in-part (CIP) of prior application No.: _____

Which is a:

Continuation Divisional Continuation-in-part (CIP) of prior application No.: _____

Which is a:

Continuation Divisional Continuation-in-part (CIP) of prior application No.: _____

Enclosed are:

Application Elements

1. Filing fee as calculated and transmitted as described below
2. Specification having 36 pages and including the following:
 - a. Descriptive Title of the Invention
 - b. Cross References to Related Applications (*if applicable*)
 - c. Statement Regarding Federally-sponsored Research/Development (*if applicable*)
 - d. Reference to Microfiche Appendix (*if applicable*)
 - e. Background of the Invention
 - f. Brief Summary of the Invention
 - g. Brief Description of the Drawings (*if drawings filed*)
 - h. Detailed Description
 - i. Claim(s) as Classified Below
 - j. Abstract of the Disclosure

UTILITY PATENT APPLICATION TRANSMITTAL (Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
PO9-99-092 (12865)

Total Pages in this Submission
59 pages

Application Elements (Continued)

3. Drawing(s) (*when necessary as prescribed by 35 USC 113*)
 - a. Formal Number of Sheets _____
 - b. Informal Number of Sheets 11
4. Oath or Declaration
 - a. Newly executed (*original or copy*) Unexecuted
 - b. Copy from a prior application (37 CFR 1.63(d)) (*for continuation/divisional application only*)
 - c. With Power of Attorney Without Power of Attorney
 - d. DELETION OF INVENTOR(S)
Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. 1.63(d)(2) and 1.33(b).
5. Incorporation By Reference (*usable if Box 4b is checked*)
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
6. Computer Program in Microfiche (*Appendix*)
7. Nucleotide and/or Amino Acid Sequence Submission (*if applicable, all must be included*)
 - a. Paper Copy
 - b. Computer Readable Copy (*identical to computer copy*)
 - c. Statement Verifying Identical Paper and Computer Readable Copy

Accompanying Application Parts

8. Assignment Papers (*cover sheet & document(s)*)
9. 37 CFR 3.73(B) Statement (*when there is an assignee*)
10. English Translation Document (*if applicable*)
11. Information Disclosure Statement/PTO-1449 Copies of IDS Citations
12. Preliminary Amendment
13. Acknowledgment postcard
14. Certificate of Mailing

First Class Express Mail (*Specify Label No.*): EJ613408276US

UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
PO9-99-092 (12865)

Total Pages in this Submission
59 pages

Accompanying Application Parts (Continued)

15. Certified Copy of Priority Document(s) *(if foreign priority is claimed)*

16. Additional Enclosures *(please identify below):*

Fee Calculation and Transmittal

CLAIMS AS FILED

For	#Filed	#Allowed	#Extra	Rate	Fee
Total Claims	17	- 20 =	0	x \$18.00	\$0.00
Indep. Claims	5	- 3 =	2	x \$78.00	\$156.00
Multiple Dependent Claims (check if applicable)					\$0.00
				BASIC FEE	\$690.00
OTHER FEE (specify purpose)					\$0.00
				TOTAL FILING FEE	\$846.00

A check in the amount of _____ to cover the filing fee is enclosed.

The Commissioner is hereby authorized to charge and credit Deposit Account No. 09-0463/IBM as described below. A duplicate copy of this sheet is enclosed.

Charge the amount of \$846.00 as filing fee.

Credit any overpayment.

Charge any additional filing fees required under 37 C.F.R. 1.16 and 1.17.

Charge the issue fee set in 37 C.F.R. 1.18 at the mailing of the Notice of Allowance, pursuant to 37 C.F.R. 1.311(b).

Signature

William A. Kinnaman, Jr., Esq.
Registration No. 27,650
IBM Corporation
Intellectual Property Law
2455 South Road, M/S P386
Poughkeepsie, New York 12601
Telephone: (914) 433-1175
Facsimile: (914) 432-9601

Dated: March 2, 2000

cc:

Docket No : **PO9-99-092**

Inventor : **Max M. Maurer**

Title : **INTELLIGENT WORKSTATION
SIMULATION-SIMULATION AT
PROTOCOL STACK LEVEL 2**

APPLICATION FOR UNITED STATES
LETTERS PATENT

"Express Mail" Mailing Label No.: **EJ613408276US**
Date of Deposit: **March 02, 2000**

I hereby certify that this paper is being deposited with the United States Postal Service as "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to: Box Patent Application, Assistant Commissioner for Patents, Washington, D.C. 20231.

Name: **Billy R. Stacy**

Signature: *Billy R. Stacy*

INTERNATIONAL BUSINESS MACHINES CORPORATION

INTELLIGENT WORKSTATION SIMULATION-SIMULATION
AT PROTOCOL STACK LEVEL 2

5

CROSS-REFERENCE TO RELATED APPLICATIONS

10

This application is related to U.S. Patent Application No. ____ (Attorney Docket No. 12864/PO999-093), entitled "INTELLIGENT WORKSTATION SIMULATION - GENERALIZED LAN FRAME GENERATION SIMULATION STRUCTURE" filed on ____, and U.S. Patent Application No. ____ (Attorney Docket No. 12866/PO999-094), entitled "INTELLIGENT WORKSTATION SIMULATION - CLIENT VIRTUALIZATION", filed on ____ both of which are incorporated herein in their entirety by reference thereto.

15

DESCRIPTION

Technical Field of the Invention

20

The present invention relates generally to a simulation method and system for testing computer networks, and more particularly to a method and system for simulating clients workstations at protocol stack level 2, the data link layer, in the layered network protocols.

25

Background of the Invention

30

Computers and computer networks currently provide important advantages to enterprises and individuals in today's society. Moreover, with the advent and ensuing popularity of the Internet and the World Wide Web ("Web"), there is a tremendous increase in volume and

PO999-092

usage of networked computer systems. Consequently, the explosive growth of computer networks has necessitated a need for larger servers to handle the network traffic. Presently, different implementation paradigms are available for handling this increase in volume. For example, some developers in the computer industry are implementing larger single servers to handle the network traffic while others are including a greater number of duplicated, relatively small servers to handle the increase in the network traffic volume.

Typically, as part of a development cycle, computer systems and software applications are tested by utilizing a simulation test tool, designed to emulate a real run time environment in order to test whether the specific computer system or application meets the various design criteria. In both of the above-described server implementations for meeting the high demand of network traffic, a complete and thorough test requires that the servicing computer, e.g., a server comprising a single machine or multiple machines, be tested at full load. Failure to execute a complete test inevitably results in a failure in the system when a full load is eventually experienced during the run-time.

Accordingly, to effectively test large and complex distributed applications and/or server applications, simulation of a realistically large client base is needed. Currently existing simulation tools provide simulation at application level. However, simulation at the application level yields an application specific simulator. Simulation at the port level is an

improvement but still fails to exercise many client-specific paths in the system under test. Thus, a general purpose simulator that also seeks to provide a high fidelity simulation, for example, a simulation at 5 level 2, the data link layer of the protocol stack is highly desirable. The data link layer is the lowest protocol stack level where each individual client has a unique client address.

10 Other simulators are currently available for testing computer networks to withstand a large volume of traffic. Figure 1 is a diagram 100 illustrating a typical prior art simulation setup. The serving system, i.e., a system under test 102 is driven by a set of simulation driving 15 systems, S 104. These driving systems 104 are controlled by a central machine, simulation controller 106. All of these machines share a common connection medium 108 such as a local area network or LAN. A plurality of such configurations may be present in order to achieve the 20 total bandwidth necessary to drive the serving system. There may be multiple controlling systems or the controlling system may have connections to multiple LANs.

25 Typically, the existing prior art simulators fall into two broad categories: 1) keystroke stuffers; and 2) protocol exercisers. A "keystroke stuffer" utilizes the actual application on the driving client. The driving vehicle is a program which retrieves stored scripts of 30 keystrokes, mouse movements, and/or other user inputs to provide input to the actual client application to drive the application. The application then directs the

traffic through a client protocol stack to the system under test 102.

Figure 2 is a diagram of the seven layer Open Systems Interconnection ("OSI") protocol stack model. Each of the layers represents a function that must be performed to effect communication between different machines. The lowest layer in the model is the physical layer 202. The functions within the physical layer 202 include setting up, maintaining, and deactivating physical circuits between systems. The most notable physical layer interfaces include IEEE 802.2 and IEEE 802.3. The next layer, i.e., layer 2, in the OSI model is the data link layer 204, which is responsible for transferring data over the physical circuits or the channel between systems. The functions of the data link layer 204 include dividing data into frames to transfer the frames to another system across the physical medium. The data link layer provides for the synchronization of data to delimit the flow of bits from the physical layer. The data link is a point-to-point link between two devices that are directly connected together.

The next layer, layer 3, is the network layer 206. The network layer 206 provides internetwork services such as the network routing and the communications between networks. The network layer 206 handles multiple point-to-point links in the case where frames are transmitted across multiple links to reach their destination.

Internet Protocol ("IP") in the Transmission Control Protocol/Internet Protocol ("TCP/IP") suite is a network layer protocol.

The next layer, layer 4, is the transport layer 208. The transport layer 208 provides end-to-end accountability of data transmitted as streams of packets. The transport layer functions include monitoring data flow to ensure proper delivery of data between source and destination.

5 It provides for error correction and for data fragmentation and reassembly. TCP is a network layer protocol.

10 The next layer up, layer 5, is the session layer 210. The session layer 210 provides for control and synchronization in exchanging data between users. For example, dialogues may be used for check and recovery of data transfer. The next layer up, layer 6, is the presentation layer 212. The presentation layer 212 functions include formatting data for display or presentation. In this layer, codes and encryption in data are interpreted and formatted for presentation. The next layer up, layer 7, is the application layer 214.

15 20 This layer is responsible for supporting end-user applications such as file transfer, electronic message exchanges and terminal session.

25 Importantly, it should be understood that all levels of the above-described OSI layers are not necessarily present in all protocol stacks. That is, the OSI protocol stack is a reference model to provide standardized logical decomposition of network into layers for communications between systems. For example. TCP/IP does not have a complete level 4 in the protocol stack because IP does not guarantee delivery of data nor will it detect missing packets, although there is a reassembly

30

function in TCP/IP that serves a portion of the function
typically found at level 4. Instead, for TCP/IP,
reliable delivery is left up to the driving application
at level 7. Thus, a keystroke stuffer is, effectively,
5 simulation above level 7, the application layer of the
protocol stack.

A "protocol exerciser" provides a somewhat more efficient
method for producing load. This method involves
10 simulation at level 7 of the protocol stack. This type
of simulator will have knowledge of a particular
application, for example, the File Transfer Protocol
("FTP"). This simulator will open ports directly with
the system under test rather than having the client
15 application do so and drive traffic to the system under
test by conforming to the higher level protocol that is
being tested, e.g., FTP, Hypertext Transfer Protocol
("HTTP"), Open Database Connectivity ("ODBC"), Simple
Mail Transfer Protocol ("SMTP"), etc. The protocol
20 exerciser provides an efficient way to generate load
because all extraneous functions of the application such
as redrawing screens, etc., do not need to be performed
at the client. A significant drawback of this type of
load generation, however, is that this simulator is
25 protocol specific. If a new protocol is developed, a new
load generator needs to be developed. For example, an
SMTP traffic generator cannot be used to test the
Internet Mail Access Protocol ("IMAP") even though both
30 protocols are mail transfer protocols. Instead, it is
necessary to develop an IMAP protocol tester.

The result of either of these methods is a series of 5
packets that pass between the system under test (102
Figure 1) and the driving system (104 Figure 1). Figure
3 is a layout of a packet for a Transmission Control
Protocol/Internet Protocol ("TCP/IP") transmission. This
packet includes a series of headers 302, 304, 306 that
have been prepended to the data 308 that the application
desires to send. Figure 3 also indicates the level 310,
10 312, 314 in the protocol stack where the data is added to
the packet and the address 316, 318, 320 that applies to
each level. The application data 308 from layer 7 is
typically passed to the transport layer protocol, for
example, the TCP. As shown in Figure 3, TCP adds a
15 header 306, and passes it to next layer protocol 312, for
example, IP. IP also adds its header 304 and passes it
to the layer 2 protocol 310 which also adds its header
302. As shown, layer 2 protocol may be Ethernet, Token
Ring, FDDI LAN. The padded data is then passed to the
physical layer where the data is converted into
20 electrical signals and transmitted to the destination
system to be received by the network interface card on
the destination system. In sum, a request from a user
traverses the OSI model until it is converted into
network traffic. Once it reaches the destination system,
25 it moves back up the OSI model so that a server
application may interpret the request.

Figure 4 shows the layout of the LAN header 400. This is
30 an IEEE 802.2 header Token ring. The shaded "Source
MACID" field 402 is the client address which is added at
level 2. MACIDs are assigned by the manufacturer of the
network interface card, and are unique to each card.

Figure 5 shows the layout of the IP header 500. The shaded "Source IP address" field 502 is the client address which is added at level 3. Figure 6 shows the layout of the TCP header. The shaded "Source Port" field 602 allows the receiving host to route the data field to a specific application running on that client. The header formats shown in Figures 4, 5, and 6 are well known to those skilled in the art of computer networking. Note that the port numbers are not client unique. That is, given a port number, it is not possible to identify the unique client to which that port number belongs. Instead, port numbers are used to distinguish between various users within the client. In contrast, the MACID and IP addresses are client unique. Given either a MACID or an IP address, a unique client can be determined. Given this overview, it is possible to examine the traffic flows resulting from simulators of the current art, to compare these flows with traffic that would result from an equal number of actual clients, and to compare and contrast how these flows would affect the system under test.

The major difference between a prior art simulation tool and real client traffic is the degree of fidelity of the simulation. Generally, the phrase "test fidelity" refers to the property that a test properly exercises code paths, for example, to closely emulate the real runtime network load in a computer network. Figure 7 shows a series of packets 700 that would originate from a prior art simulator where a single workstation was set up to simulate 100 clients. Generally, the prior art simulation method would open one hundred sessions with

the system under test. This would result in one hundred packets being sent to the system under test. However, as illustrated in Figure 7, the client specific addresses, MACID and IP address 702a..702n, 704a..704n, for each of these packets would all be identical. It is apparent to both the system under test and to anyone who might be tracing the traffic of these transactions that the test setup has a single workstation providing high volumes of traffic.

10

15

20

25

30

In contrast, Figure 8 shows one hundred packets 800 produced by a real client set. Note that each packet has unique client-specific addresses for both the MACID and the IP address 802a..802n, 804a..804n. Thus, to the system under test, this appears to be one hundred clients, each with one application running. An example of how this departure from fidelity by the prior art simulators would be visible and how it would negatively impact the test that could be achieved is explained below.

Typically, the prior art simulators asserting that each driving workstation can simulate some number of client workstations do not physically emulate multiple number of different client workstations, but one client workstation. This can be verified by using a TCP/IP provided command. TCP/IP provides a simple command: "arp -a" which displays the Address Resolution Protocol ("ARP") table. ARP determines through the use of a broadcast message packet the MACID, a physical client address, that corresponds to a specific TCP/IP address, a protocol client address. This pair of addresses is then

stored in the ARP table for future use to avoid future network traffic required for the discovery. There will be an entry in the ARP table for each workstation that communicates with a particular TCP/IP host. If this command were to be entered on the system under test, only a single entry would appear, i.e., that of the single workstation communicating with the server. In contrast, if one hundred real clients were present, the ARP table would have one hundred entries, one for each real client.

There are some cases where this distinction is not important. However, if, as an example, there were a bug in the ARP table handling code that only allowed 50 clients to be maintained, then the above test would certify this code as being good up to 100 clients while, once the code was released to customers and actually had to support 100 clients, the bug would become apparent and the code would fail. More insidiously, it would appear that the code had not been tested since the error would appear very quickly in the real environment. Thus, the example test provided by the prior art simulators provides flawed test emulation because it does not provide a realistic testing environment.

Moreover, the test is flawed because the tool used to drive the test does not provide high fidelity to the real-client environment. Specifically, the prior art simulator does not exercise all possible paths in the code implementing the system under test that the real client set exercises.

An additional disadvantage concerns the network traffic. In the above case, the ARP table is populated with the single entry for a driving workstation. This ARP table population requires a flow of data on the network. The 5 result of this flow yields the ARP table entry. In a real configuration, however, there would have to be multiple flows of network information, one to each of the real clients. Thus, in the prior art simulated environment, ARP would be under represented in the 10 network traffic flows. This would be doubly true as ARP tables have a specific size and entries tend to fall out of this table as they age. In the prior art simulated scenario, this would never occur as traffic for multiple clients is flowing to the single real workstation. In a 15 real environment, this would happen all the time. Thus, once again, the network load simulated under the prior art simulation is different than it is in production. This, once again, leaves code paths that are either not tested at all or that are under tested in the prior art simulated environment and any flaws in these code paths 20 can and will cause failures in production that will not show up in test.

The scenarios where the above-described failures occur 25 can be generalized. Specifically, if the code being tested incurs any per-client costs and/or maintains state information on a per-client basis, then high fidelity testing is required if the code paths that support those per-client items are to be sufficiently tested. Figure 9 30 illustrates this concept. The prior art simulators operate in area 1 902 providing a small number of clients each of which produces a high bandwidth. Real clients,

on the other hand, operate in the mirror reverse. Real clients operate in area 2 904, i.e., a large number of clients each of which produces a relatively small bandwidth of data. Thus, while a simulator in the prior 5 art provides the total aggregate bandwidth required by the system under test, the type of this bandwidth is incorrect in that it does not represent a correct number of clients. Testing which has this flaw will show the characteristic that in-lab testing will complete without 10 error but, upon experiencing load in a production environment, the server will fail.

Generally, the above described behavior is prominent in the services that are provided on the Internet today. 15 Referring back to Figure 9, if the number of clients is small, areas 1 902 and 2 904 overlap as shown at 906. Thus, either sort of test may be sufficient. However, as the total bandwidth grows, it is extremely important to ensure that the testing adjusts to the growth correctly. 20 As an example, it is clear that if a company production level doubles in size it is unrealistic to expect that each worker suddenly doubles their workload. Instead, it is more reasonable to expect that the number of workers doubles. Thus, growth typically occurs along the "X" 25 axis 908 of the graph in Figure 9 rather than along the "Y" axis 910. That is, as Internet usage grows, areas 1 902 and 2 904 become increasingly distant from one another. This distance represents potential failures that cannot be detected by the existing prior art 30 simulators.

The prior art simulators generally run on the machine types that they are simulating. i.e. if an Intel 80x86 machine is being simulated, the simulation generally runs on an 80x86 machine. If an IBM S/390 machine is being simulated, the simulator generally runs on a S/390 machine. This is true because current art simulators depend upon either the applications being run in the simulated platform and/or on the client protocol stack of the simulated platform.

10

Considering the foregoing, it is highly desirable to have a network simulator that is able to emulate more realistic client/server network traffic and load for testing a system under test. Particularly, it is highly desirable to have a network simulation system that is able to emulate realistically large number of clients. It is also highly desirable to have this simulation system to run on any computer system, independent of specific architecture.

15

SUMMARY OF THE INVENTION

20

The invention described herein is a method for providing high fidelity simulation of clients generating traffic to and accepting traffic from a system under test. Both generating and accepting traffic occur at level 2 of the protocol stack. Thus, a simulator utilizing the method of the present invention directly connects to the physical device that communicates with the system under test. For example, if the simulator is running on an intelligent work station, then this device is a LAN card. If the simulator is running on a mainframe, the device is

5 a communications device. If the simulator is running on the system under test, it is possible to deliver the traffic to the native protocol stack of the system under test via an internal input/output ("I/O") buffer using an I/O simulation techniques.

10 In addition to managing the physical transmissions of data, the simulator of the present invention also provides for generation of complete data packets that are client specific. Moreover, in generating complete data packets, the present invention manages the client specific data areas in the packets. The client specific data areas include application specific content data. For example, the simulator of the present invention inserts client specific addresses at levels 2 and 3 of the protocol stack and manages port and/or session numbers at level 5. Furthermore, the simulator of the present invention handles the traffic flow that is needed to maintain the conversation, for example, pacing, window control, and ARP replies, because simply "tracing" and "replaying" frames is not sufficient to emulate a real run time environment.

15 Further yet, the present invention provides for customization of data content on a per-client basis to maintain high fidelity in simulation. A simulator needs to be able to incorporate information returned from the system under test into subsequent frames. For example, tree identifiers ("TIDs") and file identifiers ("FIDs") typically identify pieces of information that the Network 20 Basic Input Output System ("NETBIOS") uses to associate open files in the Server Message Block ("SMB") protocol.

These pieces of information are client specific and will change from run to run. Thus, in simulation, a simple replay of previously observed traffic will not suffice to emulate real run time environment. These pieces of information are used in subsequent client requests.

5 Therefore, the simulator needs to capture these pieces of information from a current request and then be able to substitute this information into future requests. The simulator of the present invention is enabled to capture 10 the information from a current request for further user.

The present invention provides for simulated clients which are able to emulate the real run time environment as much as possible by having the client traverse various 15 code paths which may occur during the run time. For example, the present invention provides for realistic simulation of clients where the clients may time out due to inactivity or idling. For example, in real run time environment, many clients idle for a duration longer than 20

a predetermined amount of time. When such idling is detected, network servers typically designate these client as having timed out, and automatically terminate the network connection with the client. Such timing out, in turn may trigger ARP table timeout algorithms.

25 However, in the prior art simulators, the simulated clients tend to be constantly active. As a result of this heavy activity, the time out feature may never get invoked. To overcome such shortcomings of the prior art simulators, the present invention provides simulated 30 clients that can also emulate code paths of inactive clients.

In accordance with the goals to maintain high fidelity in simulating intelligent workstations, there is provided a simulator emulating clients at the lowest possible level where the client has a unique identity. By enabling 5 simulation of clients each having a unique identity, it is possible to provide a simulation that is indistinguishable from real client traffic. With this capability, it is possible to construct tests that provide complete coverage of the system under test both 10 in the software and in the hardware. Moreover, given the low level of the stimulation, tests can be constructed that inject errors at any level of the protocol stack.

15 Further features and advantages of the present invention as well as the structure and operation of various embodiments of the present invention are described in detail below with reference to the accompanying drawings. In the drawings, like reference numbers indicate identical or functionally similar elements.

20

BRIEF DESCRIPTION OF THE DRAWINGS

25 Preferred embodiments of the present invention will now be described, by way of example only, with reference to the accompanying drawings in which:

Figure 1 is a diagram of a typical simulator in the prior art;

30 Figure 2 is a diagram of the seven layer OSI protocol stack showing the basic functions of each layer;

Figure 3 is the layout of a TCP/IP packet;

Figure 4 shows the layout of the LAN header, 802.2 Token ring;

5 Figure 5 shows the layout of the IP header;

Figure 6 shows the layout of the TCP header;

10 Figure 7 shows packets from 100 simulated clients produced by a simulator of the current art;

Figure 8 shows packets from 100 real clients or from 100 clients simulated by this invention;

15 Figure 9 shows a graph of two distinct methods of operation for reaching a desired aggregate bandwidth at the system under test;

20 Figure 10 illustrates a simulator of the present invention that virtualizes clients at level 2 in the protocol stack;

Figure 11 is a diagram illustrating the simulator of present invention;

25 Figure 12 shows a typical application serving locally attached clients; and

Figure 13 shows a typical application serving remotely attached clients.

Figure 14 shows how a simulator of the present invention maintains independent states for simulated clients.

DETAILED DESCRIPTION OF THE
PREFERRED EMBODIMENT OF THE INVENTION

The present invention is directed to one component of a generalized client/server network simulation tool that simulates network traffic that is indistinguishable from the traffic that results from configuring a like number of real clients. The simulation tool may be employed to test host, for example, S/390 based network serving software and hardware. However, because of the generalized nature of the simulation, the scope of this tool is not limited to host testing nor to software testing. Moreover, the simulation tool of the present invention may be utilized to test the network fabric, i.e., the physical layer of the protocol stack.

To accomplish the foregoing high level of simulation fidelity, that is, a simulation closely resembling real client network paths, the network simulation tool includes: 1) simulation at level 2 in the protocol stack; 2) simulation that builds complete LAN frames; and 3) simulation that is enabled to insert and retrieve traffic for multiple clients. Each of the three components will be described briefly hereinbelow. The present invention is directed to the first component of the simulation tool, that is, the simulation at level 2 in the protocol stack.

1) Simulation at level 2 in the protocol stack

In a typical network computing environment, each client is equipped with an identifier unique to the client. This identity is unique down through level 2, the media access control or MAC layer, a part of the data link 5 layer. Specifically, each client has a unique MACID, the address that the client's LAN card recognizes in order to capture physical traffic from the LAN. To correctly simulate the real computer networking environment, the present invention provides simulated clients having 10 unique identities at level 2.

Alternatively, each client also has a unique identity at level 3, the network layer of the protocol stack. Specifically, for TCP/IP, this is the IP address for the 15 client. Furthermore, some applications tend to receive traffic at different levels depending upon their scope. For instance, Lotus Notes, which primarily serves users that are local to the enterprise and, thus, local to the LAN, tends to receive traffic at level 2. Each LAN frame 20 will have a unique MACID and a unique TCP/IP address that corresponds to the MACID. Other applications, such as Web servers, primarily serve users that are foreign to the enterprise, and thus, not attached to the LAN but who have come into the LAN through a router attached to a 25 wide area network. These servers tend to receive traffic at level 3, i.e., each LAN frame has a unique TCP/IP address that corresponds to the client but the MACID is the same in all of the frames, the MACID of the router.

30 To provide high fidelity simulation, this component of the simulation tool of the present invention provides a simulator that operates at level 2. This simulates

locally attached clients. By removing some of this generality, i.e., by setting duplicate MACIDs, simulation at level 3 can be achieved. This simulates remotely attached clients.

5

2) Simulation generating complete LAN frames

This component relates to the generation of LAN frames for the simulated clients. Level 2 is where the complete LAN frame is built just prior to transmission on the physical LAN which is level 1. Because the client identity reaches level 2 in the protocol stack, any simulator that demonstrates a high degree of fidelity must be able to generate complete LAN frames. That is, knowledge of which specific client is being simulated must pervade the protocol stack through level 2, the building of the LAN frame. A serving entity is a finite state machine, i.e., a complete transaction usually takes more than one communication. For example, in a home shopping transaction, multiple communications related to a purchase takes place between a client and the server. Building the LAN frames at level 2 enables simulation of actions of the application that would be running on the simulated client maintaining the client portion of this finite state machine. Furthermore, building LAN frames at level 2 allows retrieval and manipulation of the contents of each reply from the server as state information, thereby enabling full checking of responses from the host.

30

3) Simulation that inserts and retrieves traffic for multiple clients

Once the LAN frames are generated as described above, the frames need to be inserted on the LAN. That is, the LAN frames built at level 2 now enters level 1, the physical protocol stack for transmission. Since the LAN cards
5 generally do not check to ensure that the addresses supplied in LAN frames that are to be inserted onto the LAN are correct, it is possible to generate traffic for multiple clients and send it out. At other levels, such insertion is called "spoofing", i.e., pretending to be
10 someone else. However, the real problem arises when the server returns traffic to the simulated client. The LAN card only recognizes traffic specifically for the client identity that has been specified to the card. Thus, in
15 normal circumstances, traffic for multiple clients goes out but cannot come back in.

To circumvent this problem, this component of the simulation tool provides novel methods to be employed. These methods may be architectural specific, for example,
20 specific to token ring, or Ethernet. Moreover, they also may be generalized to cover all LAN types.

In multi-cast architecture, a LAN card may receive traffic for multiple destinations. However, multi-cast architecture does not allow individual clients to be simulated, i.e., a multicast address represents a class
25 of clients, and not a specific client in the class. Although there is a facility known as "virtual IP address" or VIP that allows one real client to have multiple identities, this facility functions at level 3, and therefore does not allow the full simulation fidelity
30 that is required.

The first component of the network simulation tool provides for the simulation at level 2, the data link layer in the protocol stack. Briefly, the data link layer defines the rules for sending and receiving information across a physical connection between two systems. Functions of the data link layer include dividing the data from the upper network layer into frames and sending the frames across the physical link. Data is framed so it can be transmitted one frame at a time. The data link layer in the receiving system will accept those frames, passing them to higher levels of the protocols. Thus, the data link layer is responsible for the transfer of data over a physical channel.

Because level 2 defines the software/hardware boundary in any given protocol stack, the simulation tool of the present invention does not require the presence of any specific protocols. Moreover, the present invention does not specify that the protocol stack simulated be TCP/IP. Thus, the present invention may encompass any protocol stack that is in existence today or any protocol stack that may be invented in the future, affording tremendous degree of extensibility which can cover new developments in the field of inter-system communications and networking. Additionally, the present invention is independent from the specification of the level 1 protocol, the physical layer, which is not a part of this invention. Consequently, the present invention can be applied to any communications type including local area network ("LAN"), wide area network ("WAN"), or any new communications protocol or architecture that may emerge.

further providing the extensibility of the present invention to any communications method.

Figure 10 illustrates a simulator 1006 that virtualizes clients at level 2 204 in the protocol stack. The mechanism of the simulator 1006 encompasses all of the functions of the protocol stack and produces fully formed packets which are delivered to the communications medium 1008. These packets then flow normally through the protocol stack on the system under test 1004 and replies are returned to the simulator 1006. The simulator 1006 captures and interprets these replies and takes appropriate actions. For example, opening a TCP/IP socket involves three packet transmissions, client to server, server to client and client to host. The initial client request flows to the server, the server sends information to the client, which is captured by the client and used to formulate a final acknowledgment from the client to the server. Because the simulator 1006 has control through level 2 (204) of the stack, the client specific addresses that are set at levels 2 and 3 (204 and 206) are under the control of the simulator 1006. In this way, the virtualization of the clients is complete and the resulting packets are indistinguishable from packets flowing from a number of distinct clients.

Figure 11 shows the physical configuration 1100 of the simulation system of the present invention both as it actually exists, shown by the solid lines 1102, and as it appears to the system under test 1004 shown by the dashed lines 1104. Because the traffic flowing is indistinguishable from traffic that is generated by

multiple clients and because the only knowledge that the system under test 1004 has of clients is their traffic, the testing system appears to the system under test 1004 to be composed of multiple real clients 1106. In fact, 5 these clients 1106 are virtualized.

By providing the ability to simulate at level 2 of a protocol stack, the present invention also enables the ability to simulate at higher levels if desirable. Most 10 significantly, being able to simulate at level 2 allows simulation at level 3. Figures 12 and 13 illustrate the difference between locally attached and remotely attached workstations. Figure 12 shows a typical application 1204 serving locally attached clients 1206. Figure 13 shows a 15 typical application 1304 serving remotely attached clients 1306. Different serving applications have different client sets that are predominantly one of these two types. As an example, a mail serving application 1204 will mostly serve locally attached clients 1206, for 20 example, the employees of the enterprise hosting the mail application. On the other hand, web serving applications 1304 will mostly serve remote clients 1306 arriving from the Internet via a level 3 router 1308, 1310. Examples 25 of remote clients 1306 include users outside the company making inquiries of the company. Thus, for some applications like the mail server 1204 shown in Figure 12, simulation at level 2 may be the most desirable. For other applications like the web server 1304 shown in Figure 13, simulation at level 3 will be the norm.

30 Generally, because the simulation of the present invention is driven at level 2 of the protocol stack, it

is not necessary that the specific architecture of the
system where the simulator is running bears any
relationship to the architecture of the system under
test. Because the simulator of present invention
5 conforms to the level 2 protocol being simulated, there
is no required relationship between the architecture of
the machine driving the test and that of the system being
tested. Indeed, serving systems are unaware of and
unconcerned about the specific architecture of the
10 systems being served. Thus, an additional advantage of
the present invention is that the simulating machines can
be chosen strictly on the merits and/or requirements of
that machine regardless of the architecture of the
machine under test. This implies an ability to scale and
15 an extensibility that does not exist in simulators of the
prior art.

The present invention may be embodied in a number of
ways. In the first embodiment, the present invention may
20 be in a form of a program running on an intelligent
workstation. In the second embodiment, the present
invention may include a program running on a mainframe.
A third embodiment may include a hardware having
intelligence to sustain the per-client data items that
25 need to be inserted in the frames produced.

In the present invention all embodiments of this
invention conform to existing standards on two points.
First, all the embodiments utilize a communications
30 device that conforms to the one used by the system under
test. For example, if the system under test uses LAN
type or WAN protocols, the embodiment needs to conform to

the LAN type used or the Wan protocols used. Second, all embodiments produce frames that conform to the protocol stack used by the system under test, e.g., TCP/IP, NETBIOS, SNA, IPX, etc. Furthermore, all embodiments of this invention generate protocol/communications medium conforming data packets and insert these packets onto the communications medium. Likewise, all embodiments provide a method for retrieving packets returned from the target system and filtering these packets to the set that are targeted to the simulated clients. All embodiments further provide the ability to capture information returned from the system under test for usage in subsequent frames.

The present invention in all its embodiments provides definition of client-specific addresses so that a plurality of clients can be simulated. Each simulated client has an individual client address at each relevant level of the protocol stack. Moreover, in the present invention conversation between each simulated client and the system under test is maintained. For example, for TCP/IP, the simulator of the present invention responds to an ARP from the system under test for each simulated client, and maintains an ARP cache for each simulated client, and any other client specific cache such as a domain name server ("DNS") for each client. These caches are maintained independently for each client.

As another TCP/IP example, the present invention opens and maintains ports for individual clients independent of any other simulated clients. Thus, within the simulated protocol stack, a client identity is maintained through

the simulation, i.e., when the simulated protocol stack is running, it is doing so on behalf of a specific client. Furthermore, the state of any individual simulated client is maintained as independent state from any other simulated clients.

Importantly, the present invention does not simply simulate the functions of the protocol stack. Between the client user and the client protocol stack sits the client application. This application maintains the state functions with the system under test. For example, when retrieving a web page, the web browser opens a socket, sends the request, then receives data until the other side closes the connection. The specific format of the data sent both ways is dictated by the HTTP protocol. This is an application level protocol that is supported by the web browser. To the protocol stack per se it is just data. The simulator of the present invention maintains the state of this conversation independently for each of the simulated clients. In addition, the simulator captures information returned from the system under test, such as tokens and session keys, for use in later transmissions. This is accomplished by allowing driving scripts to parse incoming messages and store the information thus parsed in client state variables. Subsequent transmissions can include the contents of client state variables in the requests, thus sustaining the client/host state machine.

Furthermore, the simulator of the present invention provides a driver for emulating the user of the client, i.e., the user actions. This driver may be a script

having instructions of user actions static or dynamically loaded as an input to the simulator. Alternatively, this driver may be an operator at a workstation feeding user action instructions to the simulator.

5

Figure 14 shows how the simulator of the present invention maintains independent states for simulated clients. The maintenance of unique client states for each simulated client allows high fidelity simulation desired for complete client/server testing, especially when the server seeks to establish a finite state machine with the client so that the progress of a transaction between the client and the server can be maintained. For example, in a shopping session, items are selected and added to a "shopping cart". When the shopper is satisfied with the shopper's purchases, the shopper proceeds to a "checkout" function. This function knows all of the items that are in the shopping cart which allows this function to build a bill of sale. There are well established mechanisms for doing these functions.

10

Any simulator that wishes to provide high fidelity testing must provide facilities for supporting these state maintaining mechanisms. Furthermore, pieces of state information exist at all levels in the protocol stack. That is, the client MACID exists at level 2, the client IP address exists at level 3, the TCP/IP port number exists at level 4. The above-described example describes HTTP cookies which exist at level 7 in the protocol stack, i.e., the application level. All of

15

30

these pieces of information are client specific and they must be maintained separately and independently for each simulated client.

Figure 14 illustrates how this is accomplished. Various data items are defined for each simulated client. Some of these may exist in all of the clients (e.g. the IP address and the MACID) while some may only exist for certain clients, e.g., cookies for a particular transaction type. A generalized facility called a symbol table 1402 (or some similar mechanism) is provided for use by the simulator of the present invention. Client state data is maintained in the symbol table for each individual client 1402a, ..., 1402n-1, ..., 1402n. The simulator can retrieve data 1402 via the client dispatcher 1404 from the symbol table for each client 1408 and insert the data into requests, i.e., output LAN frames 1412, flowing from the simulated client to the server. Client scripts 1406 are provided to parse incoming messages, i.e., input LAN frames 1410, retrieve data items from the messages, and store them in the symbol table 1402. Therefore, data items returned by the server to indicate state (e.g., cookies) can be returned to the server in subsequent requests, thereby maintaining the aforementioned state machine.

Because of the generalized nature of this mechanism, there are a number of other uses that can be accomplished. Specifically, if tracing of a particular client is desired, it would be possible to provide this capability on a client-by-client basis. If client specific statistics, such as, response time or transaction profiling were required, these statistics could be accumulated and displayed using the same mechanism. This mechanism not only provides a high degree of fidelity to real clients, which must also maintain such state

information, but also provides a high degree of control
over the actions of the individual clients. That is, the
piece of state information that is maintained contains
the current place where a client is executing in a script
5 that represents the various user actions. Because these
states are independent of one another, each of the
simulated clients can be executing different actions,
i.e., mimicking the behavior of real clients. This
accomplishes the objective of the present invention of
10 providing a high fidelity simulation.

While the invention has been particularly shown and
described with respect to a preferred embodiment thereof,
it will be understood by those skilled in the art that
15 the foregoing and other changes in form and details may
be made therein without departing from the spirit and
scope of the invention.

CLAIMS:

Having thus described our invention, what we claim as new, and desire to secure by Letters Patent is:

1 1. A method for providing a high fidelity
2 simulation of a client/server system including a server
3 and a plurality of locally attached intelligent client
4 workstations, the method comprising:

5 simulating at level 2 of a protocol stack by
6 formulating one or more client requests to have unique
7 client identifiers incorporated at the level 2 of a
8 protocol stack before transmitting said one or more
9 client requests on a communications channel having a
10 routing access to a server for servicing said one or more
11 client requests.

1 2. A method for providing a high fidelity
2 simulation of remotely attached intelligent workstations,
3 the method comprising:

4 simulating at level 3 of a protocol stack by
5 formulating one or more client requests to have unique
6 client identifiers and a network address associated with
7 the unique client identifiers incorporated at the level 3
8 of a protocol stack before transmitting said one or more
9 client requests on a communications channel having a
10 routing access to a server for servicing said one or more
11 client requests.

1 3. The method as claimed in claim 1, wherein the
2 step of simulating at level 2 of the protocol stack
3 includes:

4 formulating data to emulate a client request to
5 submit to the server;
6 padding the formulated data with header data that
7 conforms to communications protocol used by the server
8 receiving the client request; and
9 producing one or more level 2 data frames from the
10 padded data by inserting a unique client address into the
11 padded data, the unique client address representing a
12 unique client workstation that submitted the client
13 request.

1 4. The method as claimed in claim 3, wherein the
2 method further includes:

3 placing said one or more level 2 data frames on the
4 communications channel.

1 5. The method as claimed in claim 3, wherein the
2 method further includes:

3 maintaining independent client states for each
4 client request submitted by said unique client
5 workstation.

1 6. The method as claimed in claim 5, wherein the
2 step of maintaining independent client states further
3 includes:

4 emulating a client in an idle state to trigger a
5 timeout event.

1 7. The method as claimed in claim 3, wherein the
2 method further includes:

3 maintaining independent client states for each

4 client request submitted by a plurality of said unique
5 client workstation.

1 8. The method as claimed in claim 3, the method
2 further including:
3 incorporating static instructions that emulate user
4 actions; and
5 formulating data to emulate a client request to
6 submit to a server in response to the incorporated static
7 instructions.

1 9. The method as claimed in claim 3, the method
2 further including:
3 dynamically loading instructions that emulate user
4 actions; and
5 formulating data to emulate a client request to
6 submit to a server in response to the dynamically loaded
7 instructions.

1 10. The method as claimed in claim 3, the method
2 further including:
3 receiving operator inputs at a workstation; and
4 formulating data to emulate a client request to
5 submit to a server in response to the received operator
6 inputs.

1 11. The method as claimed in claim 1, wherein the
2 communications channel includes a local area network
3 (LAN).

1 12. The method as claimed in claim 1, wherein the
2 communications channel includes a physical communications

3 device.

1 13. The method as claimed in claim 1, wherein the
2 communications channel includes an input/output buffer
3 internal to the server.

1 14. The method as claimed in claim 2, wherein the
2 step of simulating at level 3 of the protocol stack
3 includes:

4 formulating data to emulate a client request to
5 submit to the server;

6 padding the formulated data with header data that
7 conforms to communications protocol used by the server
8 receiving the client request; and

9 producing one or more level 3 data frames from the
10 padded data by inserting a unique client address and a
11 network address associated the unique client address into
12 the padded data, the unique client address representing a
13 unique client workstation that submitted the client
14 request.

1 15. A program storage device readable by machine,
2 tangibly embodying a program of instructions executable
3 by the machine to perform the method steps of providing a
4 high fidelity simulation of locally attached intelligent
5 workstations, the method steps comprising:

6 simulating at level 2 of a protocol stack by
7 formulating one or more client requests to have unique
8 client identifiers incorporated at the level 2 of a
9 protocol stack before transmitting said one or more
10 client requests on a communications channel having a

11 routing access to a server for servicing said one or more
12 client requests.

16. A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform the method steps of providing a high fidelity simulation of remotely attached intelligent workstations, the method steps comprising:

simulating at level 3 of a protocol stack by formulating one or more client requests to have unique client identifiers and a network address associated with the unique client identifiers incorporated at the level 3 of a protocol stack before transmitting said one or more client requests on a communications channel having a routing access to a server for servicing said one or more client requests.

17. An apparatus for providing a high fidelity simulation of a client/server system including a server and a plurality of locally attached intelligent client workstations, the apparatus comprising:

a server system under test receiving one or more packets of data;

a simulator that inserts a unique client identifier to said one or more packets of data at level 2 of a protocol stack, the unique client identifier representing a unique client workstation; and

a communication channel interconnecting the server system and the simulator, the simulator placing said one or more packets of data on the communication channel for transmission.

ABSTRACT OF THE INVENTION

A general purpose simulator providing a high fidelity simulation performs the simulation at level 2, the data link layer of the protocol stack level where each individual client has a unique client address. The unique client address emulated by the simulator is inserted into a data frame at the data link layer of a protocol stack before the data frame is transmitted on the communications channel. Simulation at level 3 is also performed by inserting level 3 network identifier to emulated remote multiple clients. The simulator of the present invention may be run at a workstation residing connected to a server servicing the simulated client requests. The simulator also may be run on a host computer connected to a server via a communications device. Additionally, the simulator may also be run on the same system as the server.

MAX M. MAURER

P09-99-092

1 of 11

100

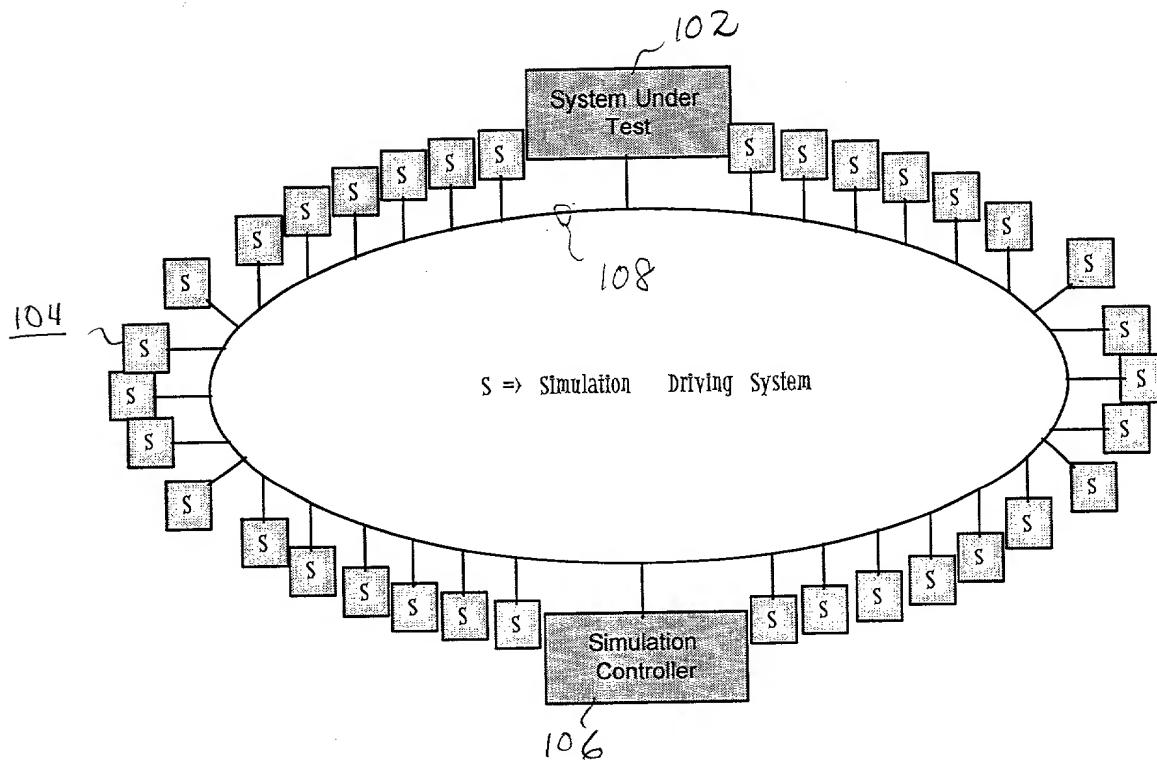


Figure 1

PRIOR ART

200

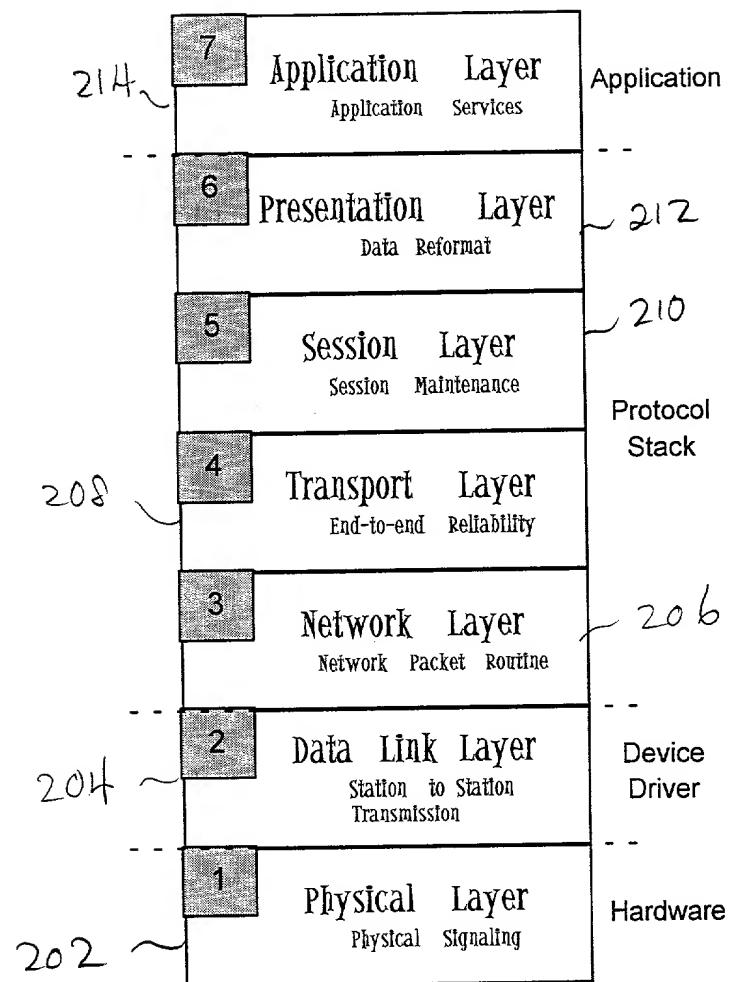


Figure 2

PRIOR ART

P09-99-092

3 of 11

300

TCP/IP Packet Elements

Data added by ...

LAN Header	IP Header	TCP Header	Application Data
------------	-----------	------------	------------------

310 312 314 322

302 304 306 308

316 317 319 320 321 323

Figure 3

Figure 3

PO9-99-092

4 of 11

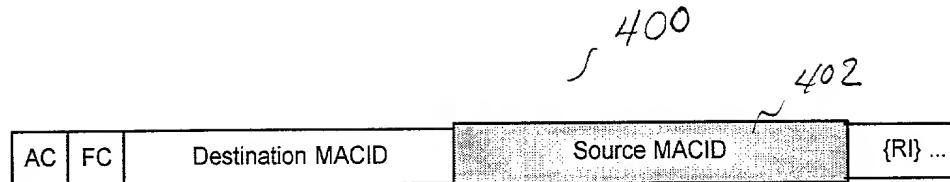


Figure 4 PRIO R ART

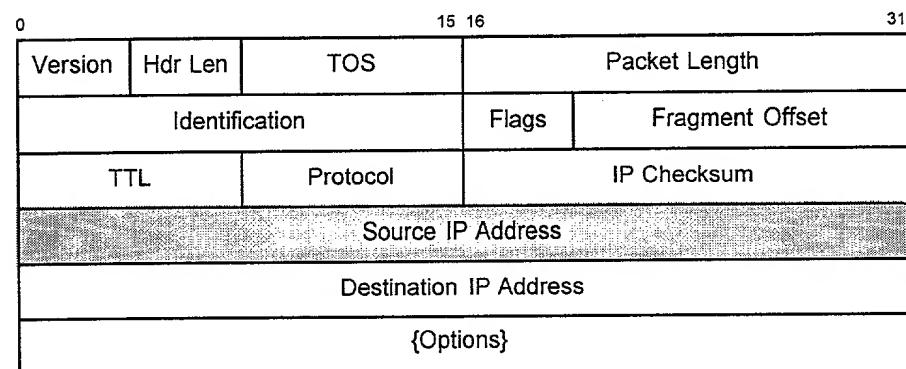


Figure 5 PRIO R ART

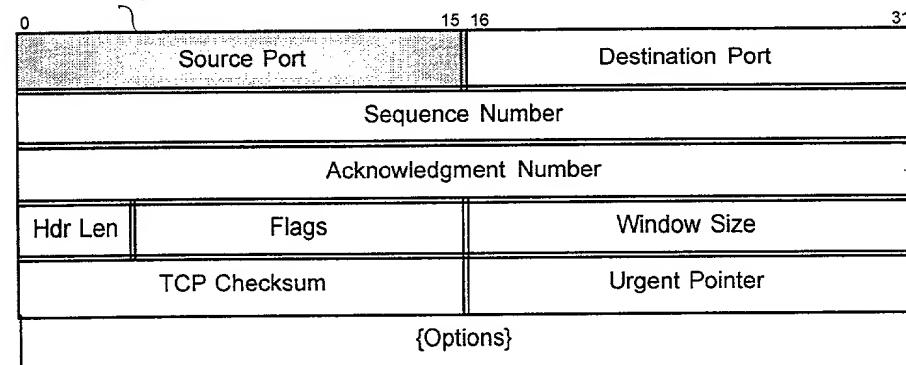


Figure 6

PRIO R ART

POG-92-092

5 of 11

700
ʃ

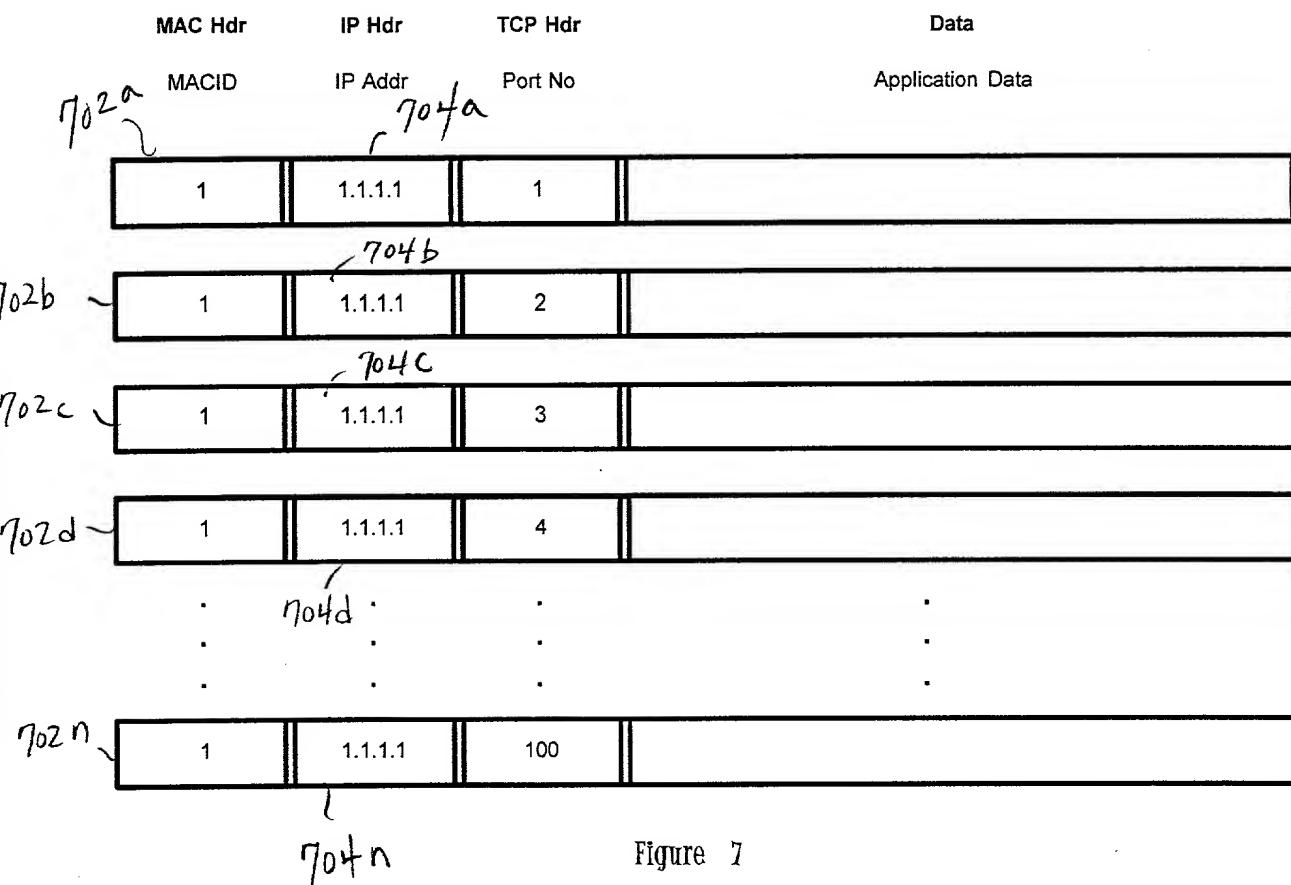


Figure 7

PRIOR ART

P09-99-092

6 of 11

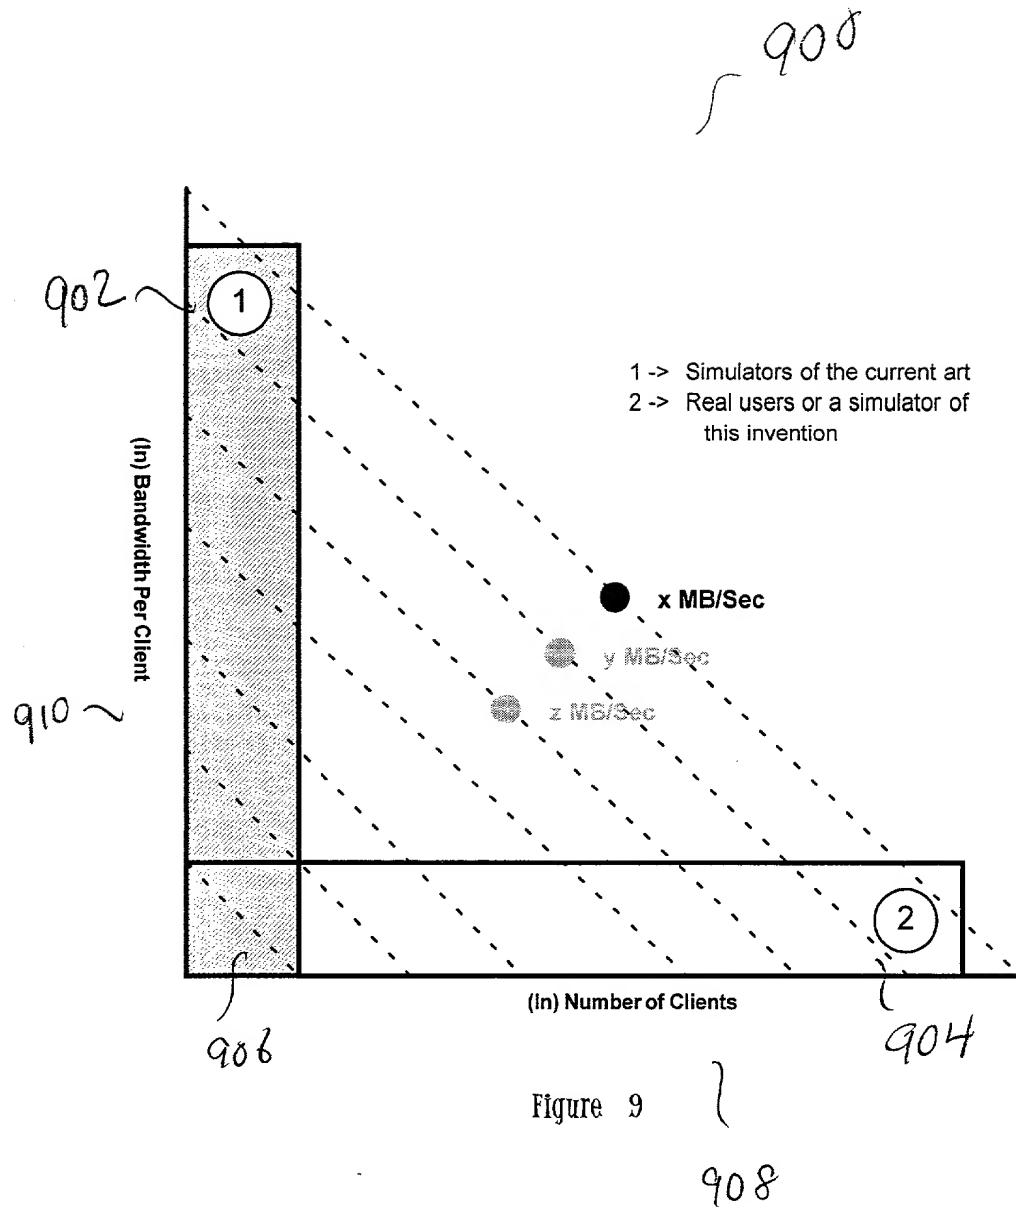
800

MAC Hdr	IP Hdr	TCP Hdr	Data
MACID	IP Addr	Port No	Application Data
1	1.1.1.1	1	
2	1.1.1.2	1	
3	1.1.1.3	1	
4	1.1.1.4	1	
.	.	.	.
.	.	.	.
802n	804n	.	.
100	1.1.1.100	1	

Figure 8

PO9-99-092

7 of 11



PO9-99-092

8 of 11

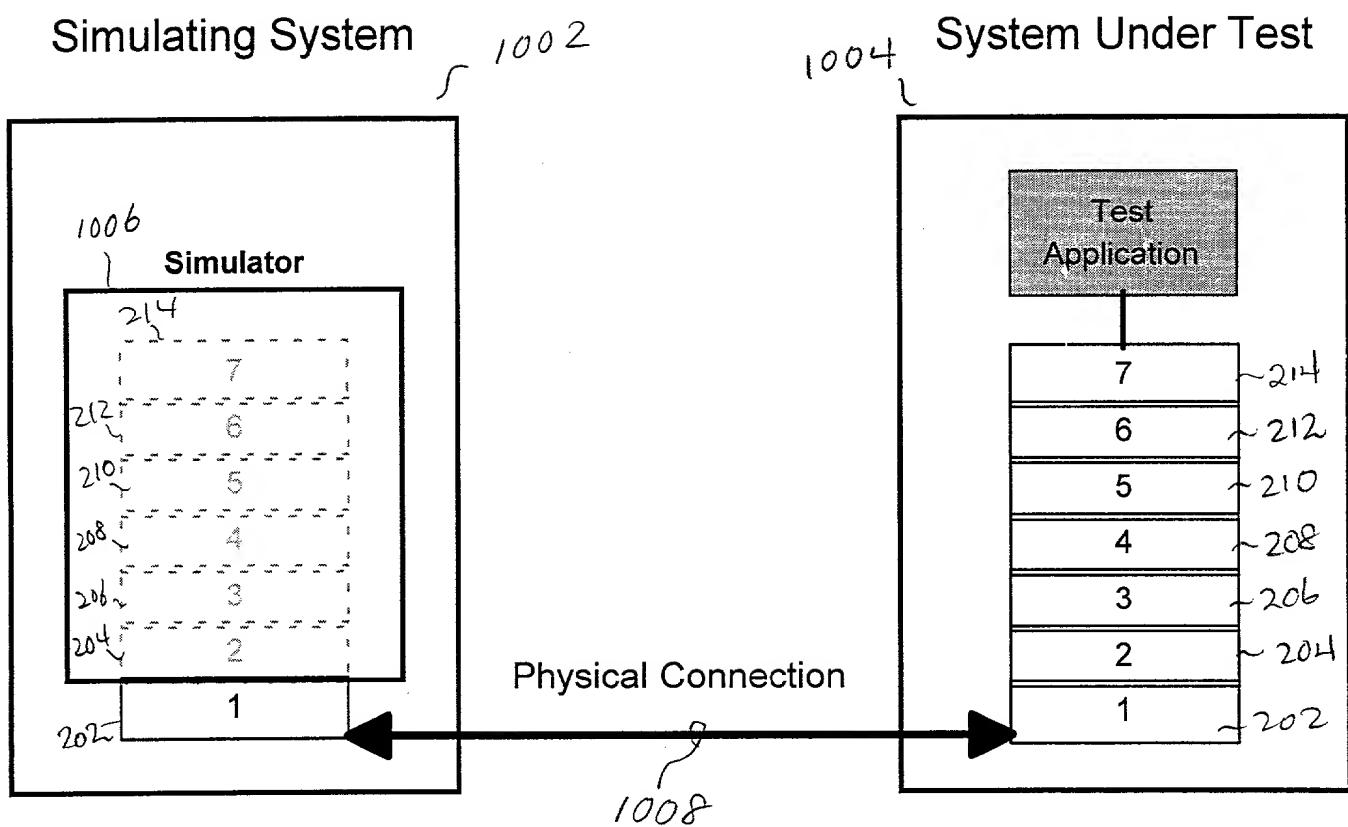


Figure 10

POG-99-092

9 of 11

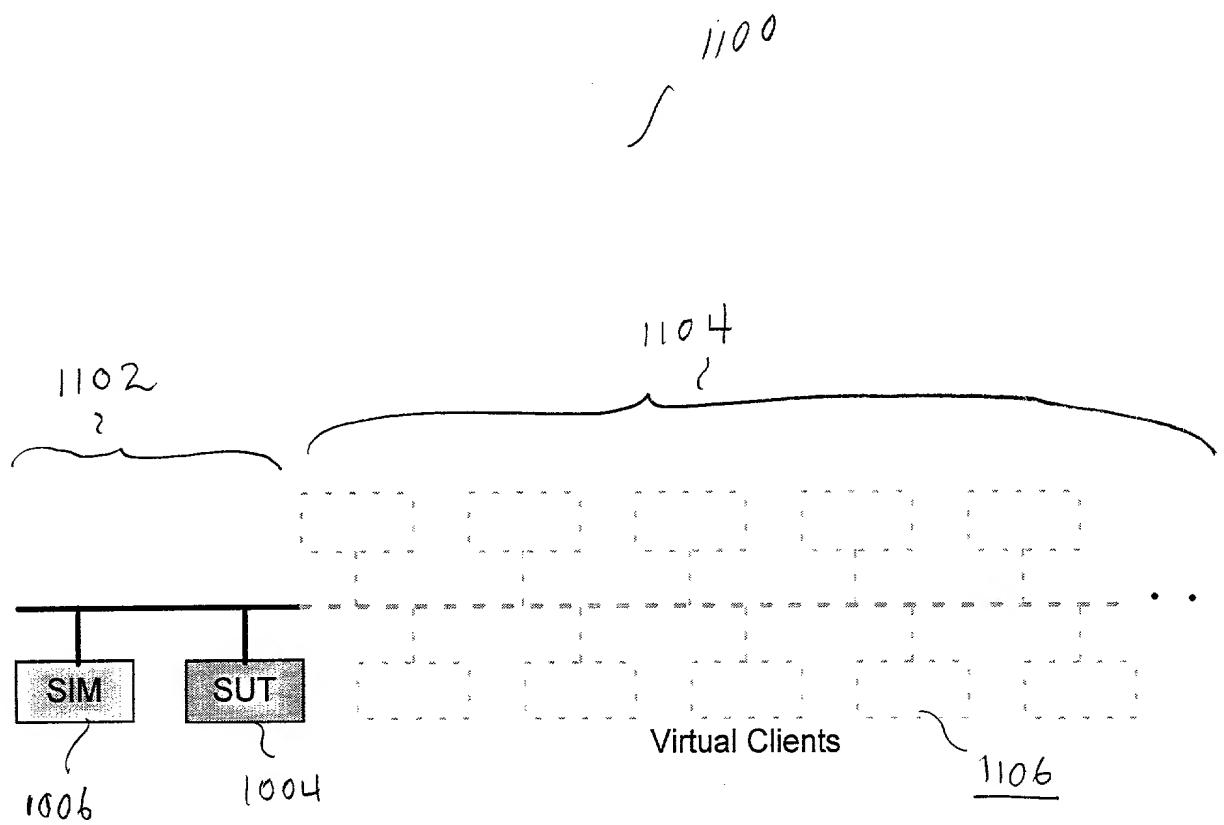


Figure 11

PO9-99-092

10 of 11

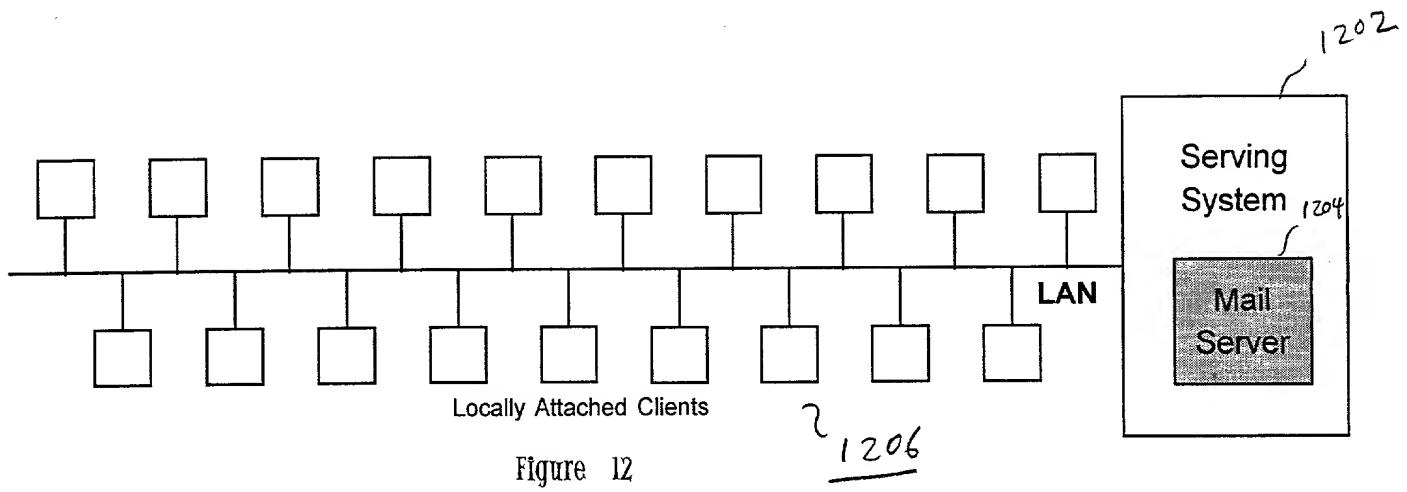


Figure 12

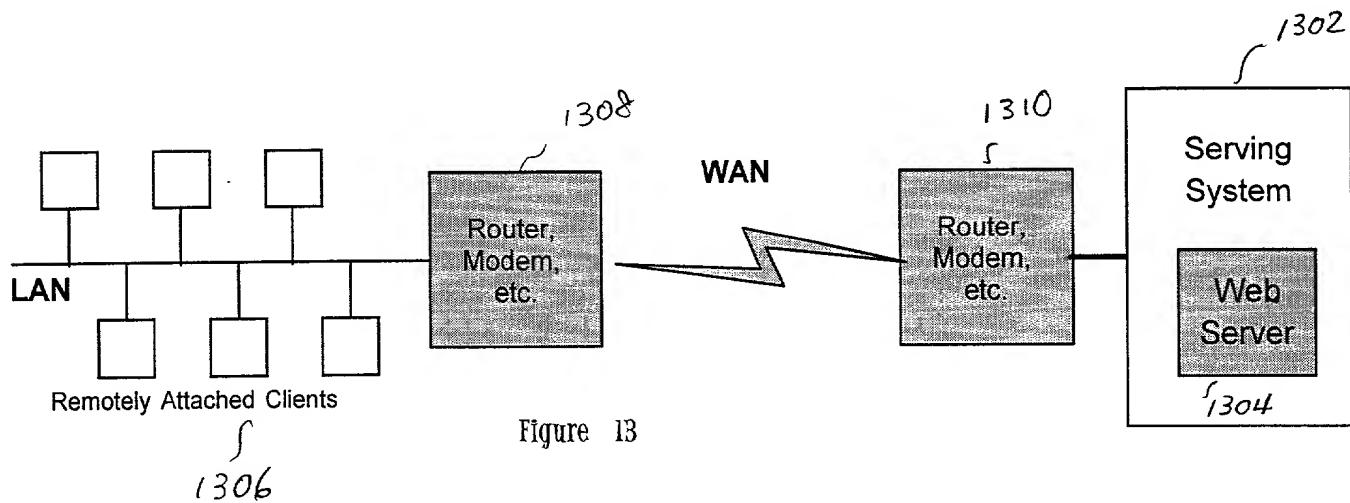


Figure 13

P09-99-092

11 of 11

~1400

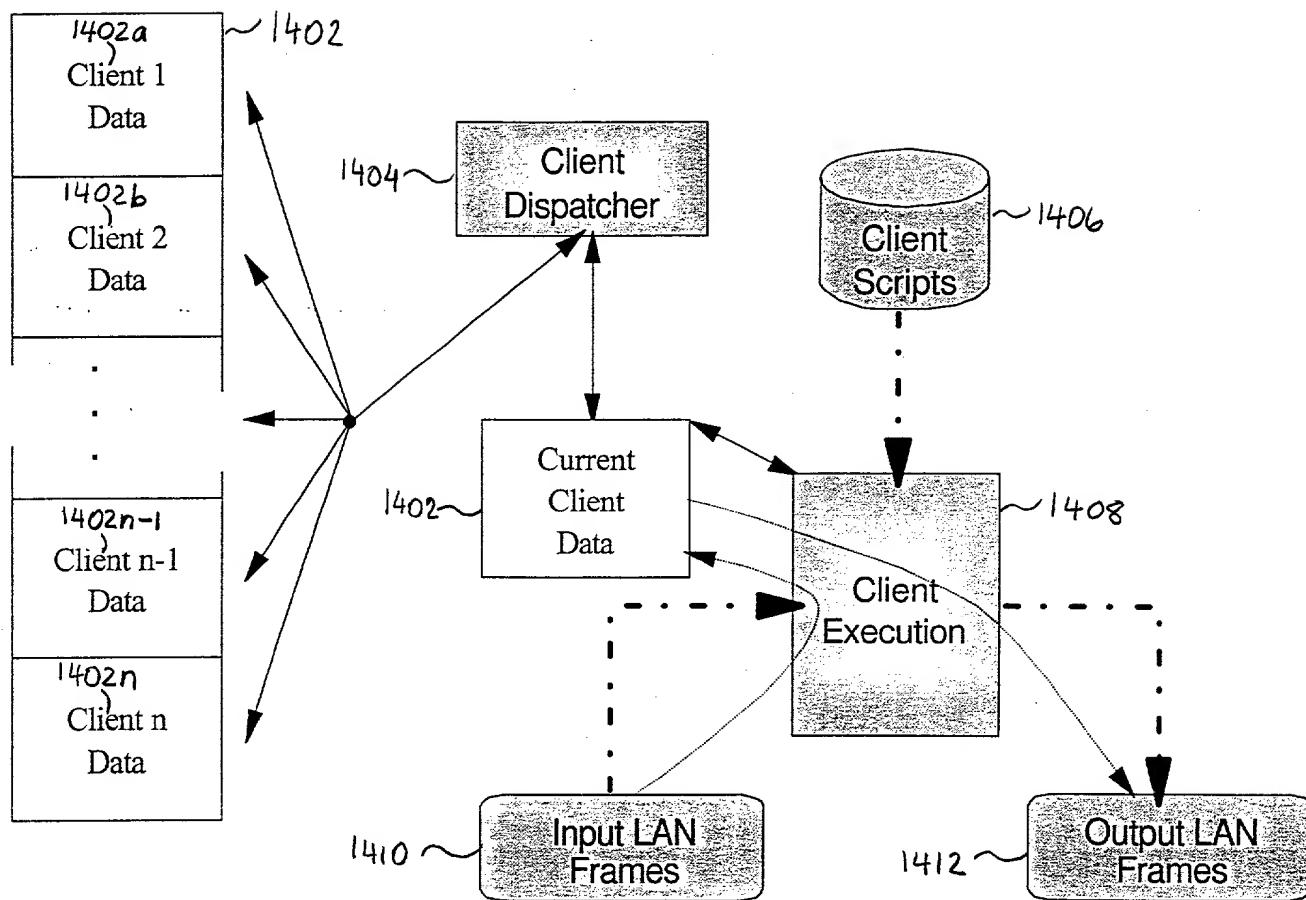


Figure 14

Declaration and Power of Attorney for Patent Application

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name; I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: **INTELLIGENT WORKSTATION SIMULATION - SIMULATION AT PROTOCOL STACK LEVEL 2**

the specification of which (check one)

X

is attached hereto.

was filed on _____ as Application Serial No. _____ and was amended on _____.

I hereby state that I have reviewed and understand the contents of the above- identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s):

Number	Country	Day/Month/Year	Priority Claimed
--------	---------	----------------	------------------

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose information material to the patentability of this application as defined in Title 37, Code of Federal Regulations, §1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

Prior U.S. Applications:

Serial No.	Filing Date	Status
------------	-------------	--------

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

As a named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith: Lynn L. Augspurger (Reg No. 24,277); Marc A. Ehrlich (Reg. No. 39,966); Floyd A. Gonzalez (Reg. No. 26,732); Lily Neff (Reg. No. 38,254); Lawrence D. Cutter (Reg. No. 28,501); William B. Porter (Reg. No. 33,135); William A. Kinnaman, Jr. (Reg. No. 27,650); Christopher A. Hughes (Reg. No. 26,914); Edward A. Pennington (Reg. No. 32,588); John E. Hoel (Reg. No. 26,279); Joseph C. Redmond, Jr. (Reg. No. 18,753); Richard L. Catania, (Reg. No. 32,608); Leopold Presser, (Reg. No. 19,827); Stephen D. Murphy (Reg. No. 22,002); Frank S. DiGilio (Reg. No. 31,346); Kenneth L. King (Reg. No. 24,223) and Edward W. Grolz (Reg. No. 33,705).

Send all correspondence to:

Richard L. Catania, Esq.
Scully, Scott, Murphy & Presser
400 Garden City Plaza
Garden City, NY 11530

Direct Telephone Calls to: Richard L. Catania., (516) 742-4343

Declaration and Power of Attorney for Patent Application

(1) Inventor: Max M. Maurer

Signature: 

3-2-00

Date

Residence: 2 Stokes Lane
Gardiner, New York 12525

Citizenship: USA

Post Office Address: same as residence